arXiv:1404.6763v2 [cs.DS] 8 May 2014

Semi-Streaming Set Cover
(Full Version)

Yuval Emek* Adi Rosén†

Abstract

This paper studies the set cover problem under the semi-streaming model. The underlying set system is formalized in terms of a hypergraph $G = (V, E)$ whose edges arrive one-by-one and the goal is to construct an edge cover $F \subseteq E$ with the objective of minimizing the cardinality (or cost in the weighted case) of F. We consider a parameterized relaxation of this problem, where given some $0 \leq \epsilon < 1$, the goal is to construct an edge $(1-\epsilon)$-cover, namely, a subset of edges incident to all but an ϵ-fraction of the vertices (or their benefit in the weighted case). The key limitation imposed on the algorithm is that its space is limited to (poly)logarithmically many bits per vertex.

Our main result is an asymptotically tight trade-off between ϵ and the approximation ratio: We design a semi-streaming algorithm that on input graph G, constructs a succinct data structure $\mathcal{D}$ such that for every $0 \leq \epsilon < 1$, an edge $(1-\epsilon)$-cover that approximates the optimal edge (1-)cover within a factor of $f(\epsilon, n)$ can be extracted from $\mathcal{D}$ (efficiently and with no additional space requirements), where

$$f(\epsilon, n) = \begin{cases} O(1/\epsilon), & \text{if } \epsilon > 1/\sqrt{n} \\ O(\sqrt{n}), & \text{otherwise} \end{cases} .$$

In particular for the traditional set cover problem we obtain an $O(\sqrt{n})$-approximation. This algorithm is proved to be best possible by establishing a family (parameterized by ϵ) of matching lower bounds.

*Technion, Israel. Email: yemek@ie.technion.ac.il.

†CNRS and Université Paris Diderot, France. Email: adiro@liafa.univ-paris-diderot.fr. Research supported in part by ANR project RDAM.

1 Introduction

Given a *set system* consisting of a *universe* of items and a collection of item sets, the goal in the *set cover* problem is to construct a minimum cardinality subcollection of sets that covers the whole universe. This problem is fundamental to combinatorial optimization with applications ranging across many different domains. It is one of the 21 problems whose NP-hardness was established by Karp in [12] and its study has led to the development of various techniques in the field of approximation algorithms (see, e.g., [21]).

In this paper, we investigate the set cover problem under the *semi-streaming* model [6], where the sets arrive one-by-one and the algorithm's space is constrained to maintaining a small number of bits per item (cf. the *set-streaming* model of [19]). In particular, we are interested in the following two research questions: (1) What is the best approximation ratio for the set cover problem under such memory constraints? (2) How does the answer to (1) change if we relax the set cover notion so that the set subcollection is required to cover only a δ-fraction of the universe?

On top of the theoretical interest in the aforementioned research questions, studying the set cover problem under the semi-streaming model is justified by several practical applications too. For example, Saha and Getoor [19] describe the setting of a web crawler that iterates a large collection of blogs, listing the topics covered by each one of them. A user interested in a certain set of topics can run a semi-streaming set cover algorithm with relatively small memory requirements to identify a subcollection of blogs that covers her desired topics.

The model. In order to fit our terminology to the graph theoretic terminology traditionally used in the semi-streaming literature (and also to ease up the presentation), we use an equivalent formulation for the set cover problem in terms of edge covers in hypergraphs: Consider some *hypergraph* $G = (V, E)$, where V is a set of n *vertices* and E is a (multi-)set of m *hyperedges* (henceforth *edges*), where each edge $e \in E$ is an arbitrary non-empty subset $e \subseteq V$. Assume hereafter that G does not admit any isolated vertices, namely, every vertex is incident to at least one edge. We say that an edge subset $F \subseteq E$ *covers* G if every vertex in V is incident to some edge in F. The goal of the *edge cover* problem is to construct a subset $F \subseteq E$ of edges that covers G, where the objective is to minimize the cardinality $|F|$.

A natural relaxation of the covering notion asks to cover some fraction of the vertices in V: Given some $0 < \delta \leq 1$, we say that an edge subset $F \subseteq E$ *δ-covers* G if at least δn vertices are incident to the edges in F, namely, $|V(F)| \geq \delta n$, where $V(F) = \{v \in V \mid \exists e \in F \text{ s.t. } v \in e\}$. Under this terminology, a cover of G is referred to as a 1-cover. This raises a bi-criteria optimization version of the set cover problem, where the goal is to construct an edge subset $F \subseteq E$ that δ-covers G with the objective of minimizing $|F|$ and maximizing δ. In this paper, we focus on approximation algorithms, where the cardinality of F is compared to that of an optimal edge (1-)cover of G.

In the *weighted* version of the edge cover problem, the hypergraph G is augmented with vertex *benefits* $b : V \to \mathbb{Q}_{>0}$ and edge *costs* $c : E \to \mathbb{Q}_{>0}$. The edge cover definition is generalized so that edge subset $F \subseteq E$ is said to *δ-cover* G if the benefit of the vertices incident to the edges in F is at least a δ-fraction of the total benefit, namely, $b(V(F)) \geq \delta \cdot b(V)$, where $b(U) = \sum_{v \in U} b(v)$ for every vertex subset $U \subseteq V$. The goal is then to construct an edge subset F that δ-covers $G = (V, E, b, c)$, where the objective is to maximize δ and minimize the cost of F, denoted $c(F) = \sum_{e \in F} c(e)$.

Under the *semi-streaming* model, the execution is partitioned into discrete time steps and the edges in E are presented one-by-one so that edge $e_t \in E$ is presented at time $t = 0, 1, \ldots, m-1$, listing all vertices $v \in e_t$;[1] in the weighted version, the cost of e_t and the benefits of the vertices it contains are also listed. The key limitation imposed on the algorithm is that its space is limited; specifically, we allow the algorithm to maintain $\log^{O(1)} |G|$ bits per vertex, where $|G|$ denotes the number of bits in the standard binary encoding of G. Each edge $e \in E$ is associated with a unique *identifier* $\mathrm{id}(e)$ of size $O(\log m)$ bits, say, the time t at which edge e_t is presented. We may sometimes use the identifier $\mathrm{id}(e)$ when we actually refer to the edge e itself, e.g., replacing $c(e)$ with $c(\mathrm{id}(e))$; our intention will be clear from the context.

In contrast to the random access memory model of computation, where given a collection $\mathcal{I}$ of identifiers, one can easily determine which vertex in V is incident to which of the edges whose identifiers are in $\mathcal{I}$ simply by examining the input, under the semi-streaming model, the collection $\mathcal{I}$ by itself typically fails to provide this information. Therefore, instead of merely returning the identifiers of some edge δ-cover, we require that the algorithm outputs a *δ-cover certificate* χ for G which is a partial function from V to $\{\mathrm{id}(e) \mid e \in E\}$ with *domain*

$$\mathrm{Dom}(\chi) = \{v \in V \mid \chi \text{ is defined over } v\}$$

and *image*

$$\mathrm{Im}(\chi) = \{\mathrm{id}(e) \mid \exists v \in \mathrm{Dom}(\chi) \text{ s.t. } \chi(v) = \mathrm{id}(e)\}$$

that satisfies (1) if $v \in \mathrm{Dom}(\chi)$ and $\chi(v) = \mathrm{id}(e)$, then $v \in e$; and (2) $b(\mathrm{Dom}(\chi)) \geq \delta \cdot b(V)$. By definition, the image of χ consists of the identifiers of the edges in some edge δ-cover F of G and the quality of the δ-cover certificate χ is thus measured in terms of $c(\mathrm{Im}(\chi)) = c(F)$.

Our contribution. Consider some unweighted hypergraph $G = (V, E)$ with optimal edge 1-cover `OPT`. We design a deterministic semi-streaming algorithm, referred to as `SSSC` (acronym of the paper's title), for the edge (δ-)cover problem that given some $0 \leq \epsilon < 1$, outputs a $(1-\epsilon)$-cover certificate χ_ϵ for G with image of cardinality $|\mathrm{Im}(\chi_\epsilon)| = O(\min\{1/\epsilon, \sqrt{n}\} \cdot |\texttt{OPT}|)$.[2] This result is extended to the weighted case, where $G = (V, E, b, c)$, showing that $c(\mathrm{Im}(\chi_\epsilon)) = O(\min\{1/\epsilon, \sqrt{n}\} \cdot$

[1] With the exception of our related work discussion, all semi-streaming algorithms in this paper make a single (one way) pass over the input hypergraph.

[2] Define $\min\{1/x, y\} = y$ when $x = 0$.

$c(\texttt{OPT})$) (see Thm. 2.2 and 2.3). In particular, for the edge (1-)cover problem, we obtain an $O(\sqrt{n})$-approximation for both the weighted and unweighted cases.

On the negative side, we prove that for every $\epsilon \geq 1/\sqrt{n}$, if a randomized semi-streaming algorithm for the set cover problem outputs a $(1-\epsilon)$-cover certificate χ for G, then it cannot guarantee that $\mathbb{E}[|\mathrm{Im}(\chi)|] = o(|\texttt{OPT}|/\epsilon)$ (see Thm. 3.1). This demonstrates that the approximation guarantee of our algorithm is asymptotically optimal for the whole range of parameter $0 \leq \epsilon < 1$ even for randomized algorithms.

Notice that `SSSC` has the attractive feature that the (near-linear size) data structure $\mathcal{D}$ it maintains is oblivious to the parameter ϵ. That is, the algorithm processes the stream of edges with no knowledge of ϵ, generating the data structure $\mathcal{D}$, and the promised $(1-\epsilon)$-cover certificate χ_ϵ can be efficiently extracted from $\mathcal{D}$ (with no additional space requirements) for every $0 \leq \epsilon < 1$ (in fact several such covers for different values of ϵ can be extracted). From a bi-criteria optimization perspective, our lower bound implies that the parameterized collection $\{\chi_\epsilon\}_{0\leq\epsilon<1}$ encoded in $\mathcal{D}$ is an (asymptotically) optimal solution frontier (cf. Pareto optimality).

Using a simple adjustment of the randomized rounding technique for set cover (see, e.g., [21]), it is not difficult to show that a basic feasible solution to the linear program relaxation $\mathcal{P}$ of a given set cover instance also serves as a compact data structure from which a $(1-\epsilon)$-cover certificate χ_ϵ can be extracted for every $0 \leq \epsilon < 1$. In fact, the approximation ratio obtained this way is better than ours, namely, $O(\log(1/\epsilon))$. However, our lower bound shows that this approach cannot be applied — and in passing, that $\mathcal{P}$ cannot be solved — under the semi-streaming model.

Can our tight lower bound be an artifact of the requirement that the algorithm outputs a cover *certificate*? We nearly eliminate this possibility by proving that for every constant $c > 0$ and for every $\epsilon \geq n^{-1/2+c}$, even if the randomized algorithm only guarantees an "uncertified" output, i.e., only the identifiers of the edges in some edge $(1-\epsilon)$-cover F of G are returned, then the cardinality of F must still be large, specifically, $|F| = \Omega\left(\frac{\log\log n}{\log n} \cdot |\texttt{OPT}|/\epsilon\right)$, where `OPT` in this case is proportional to $\epsilon^2 n$ (see Thm. 3.2).[3]

Related work. The work most closely related to the present paper is probably the one presented in Saha and Getoor's paper [19] that also considers the set cover problem under the semi-streaming model (referred to as *set-streaming* in [19]) formulated as the edge cover problem in hypergraphs. Saha and Getoor design a 4-approximation semi-streaming algorithm for the *maximum coverage* problem that given a hypergraph $G = (V, E)$ and a parameter k, looks for k edges that cover as

[3] By using a reduction from the index function studied in communication complexity [15], one can show that there does not exist a semi-streaming algorithm that distinguishes between hypergraphs admitting a constant size edge cover and hypergraphs that cannot be covered by less than n^α edges for any constant $0 < \alpha < 1/2$. This lower bound is more attractive in the sense that it applies already to the decision version of the set cover problem however, to the best of our understanding, in contrast to the constructions of the present paper, this result cannot be generalized to $(1-\epsilon)$-covers for values of $\epsilon \gg 1/\sqrt{n}$.

many vertices as possible. Based on that, they observe that an $O(\log n)$-approximation for the optimal set cover can be obtained in $O(\log n)$ passes over the input (this can be achieved based on our semi-streaming algorithm as well). Using the terminology of the present paper, Saha and Getoor's maximum coverage algorithm is very efficient for obtaining edge $(1-\epsilon)$-covers as long as ϵ is large, but it does not provide any (single pass) guarantees for $\epsilon < 3/4$. In contrast, our algorithm has asymptotically optimal (single pass) guarantees for any $0 \leq \epsilon < 1$. Another paper that considers semi-streaming algorithms in hypergraphs is that of Halldórsson et al. [10] that studies the independent set problem.

The semi-streaming model was introduced by Feigenbaum et al. [6] for graph theoretic problems, where the edges of an n vertex input graph arrive sequentially and the algorithm is allowed to maintain only $\log^{O(1)} n$ bits of memory per vertex. Since the number of bits required to encode an n vertex graph is $n^{O(1)}$, the space-per-vertex bound used in the present paper can be viewed as a generalization of that of Feigenbaum et al. from graphs to hypergraphs. In any case, concerns regarding the comparison between the space bound used in the present paper and that of [6] can be lifted by restricting attention to hypergraphs with $m \leq 2^{\log^{O(1)} n}$ edges (refer to Sec. 2 for a further discussion of the space bounds of our algorithm).

Various graph theoretic problems have been treated under the semi-streaming model. These include matching [17, 5, 14], diameter and shortest path [6, 7], min-cut and sparsification [1, 13], graph spanners [7], and independent set [10, 4].

Several variants of the set cover problem, all different than the problem studied in the present paper, have been investigated under the model of online computation. Alon et al. [2] focus on the online problem in which some master set system is known in advance and an unknown subset of its items arrive online; the goal is to cover the arriving items, minimizing the number of sets used for that purpose. Another online variant of the set cover problem is studied by Fraigniaud et al. [8], where the sets arrive online, but not all items have to be covered. Here, each item is associated with a penalty and the cost of the algorithm is the sum of the total cost of the sets chosen for the partial cover and the total penalty of the uncovered items.

Note that under the online computation model, there is a trivial linear lower bound for the problem studied in the present paper if preemption is not allowed. If preemption is allowed, then the problem becomes interesting only under a slightly stronger definition for the competitive ratio: The performance of the algorithm is measured via the maximum over time t of the ratio $\texttt{ALG}_t/\texttt{OPT}_t$, where $\texttt{OPT}_t$ is the cost of an optimal set cover for the set system presented up to time t and $\texttt{ALG}_t$ is the cost of the set cover maintained by the algorithm for that set system. The set cover algorithm presented in the present paper is, in fact, also an online algorithm for this problem with competitive ratio $O(\sqrt{n})$. The lower bound(s) established in the present paper can be slightly modified to show that this is optimal.

Closely related to our notion of cover certificate is the *universal set cover* problem [11, 9], where

given a set system, the goal is to construct a mapping f from the items to the sets containing them so that for every item subset X, the cost of the image of X under f is as close as possible to the cost of a minimum set cover for X. This problem resembles our guarantee that the promised $(1-\epsilon)$-cover certificate can be extracted from the data structure for every ϵ however, it is much stronger in the sense that it guarantees a small cover for every item subset, rather than the existence of a "good" item subset for every ϵ. To the best of our knowledge, the universal set cover problem has not been studied under the semi-streaming model.

Techniques' overview. The main procedure of our algorithm `SSSC` (referred to as `COVER`) maintains for each vertex $v \in V$, a variable $\text{eff}(v)$. This variable captures the ratio of the benefit of the last *effective* subset $T \subseteq e_t$ that covered v to the cost of e_t, where subset $T \subseteq e_t$ is said to be effective if $b(T)/c(e_t) \geq 2 \cdot \text{eff}(u)$ for every $u \in T$. This means, in particular, that the variable $\text{eff}(v)$ doubles with every update. (Note that `COVER` actually maintains the logarithm of this $\text{eff}(v)$ variable for each vertex v, but the main idea is the same.) By picking the effective subset $T \subseteq e_t$ that maximizes $b(T)$, we ensure that the collection of vertices $v \in V$ admitting high values of $\text{eff}(v)$ satisfies some desirable properties. Specifically, a careful analysis shows that upon termination of the input stream, there exists some threshold ρ such that the total benefit of vertices $v \in V$ with $\text{eff}(v) \leq \rho$ is at most $\epsilon \cdot b(V)$, whereas the total cost of the edges corresponding to the effective subsets of the vertices $v \in V$ with $\text{eff}(v) > \rho$ is $O(c(\texttt{OPT})/\epsilon)$. Invoking procedure `COVER` on a hypergraph with the same edge costs and uniform vertex benefits (in parallel to the invocation of `COVER` on the original input hypergraph) enables us to produce an edge 1-cover that $O(\sqrt{n})$-approximates $c(\texttt{OPT})$.

The bad hypergraphs that lie at the heart of our lower bound are constructed based on an *affine plane* $\mathcal{A} = (P, L)$ with q^2 points and $q(q+1)$ lines (see, e.g., [16]) by randomly partitioning each line in L into two edges (more edges in the "uncertified" version of the lower bound). After presenting the two edges corresponding to all lines in L, we present one additional edge e^* that contains the points of all but $r \approx \epsilon q$ random lines from some random angle A_i of $\mathcal{A}$. An optimal edge cover consists of the edge e^* and the $2r = O(\epsilon q)$ edges corresponding to the r lines missing from e^*. Using careful information theoretic arguments, we show that any low space deterministic algorithm must use many lines from angles other than A_i to construct a $(1-\epsilon)$-cover F. The properties of affine planes guarantee that the expected cardinality of F is $\Omega(q)$. By Yao's principal, our lower bound is translated from deterministic algorithms to randomized ones.

2 A semi-streaming algorithm

Our goal in this section is to design a semi-streaming algorithm for the edge (δ-)cover problem in hypergraphs. The algorithm, referred to as `SSSC`, is presented in Sec. 2.1 and its approximation ratio is analyzed in Sec. 2.2. For the sake of simplicity, we first assume that all numerical values

(vertex benefits and edge costs) are encoded using $O(\log n)$ bits. Under this assumption, the space bounds of SSSC are quite trivial and the analysis in Sec. 2.2 yields Theorem 2.1.

Theorem 2.1. *On a weighted input hypergraph $G = (V, E, b, c)$ with numerical values encoded using $O(\log n)$ bits, our algorithm uses $O(n \log(n + m))$ space, processes each input edge $e_t \in E$ in $O(|e_t| \log |e_t|)$ time, and produces a data structure $\mathcal{D}$ with the following guarantee: For every $0 \leq \epsilon < 1$, a $(1 - \epsilon)$-cover certificate χ_ϵ for G such that*

$$c(\mathrm{Im}(\chi_\epsilon)) = O\left(\min\left\{1/\epsilon, \sqrt{n}\right\} \cdot c(\texttt{OPT})\right)$$

can be extracted from $\mathcal{D}$ in time $O(n \log n)$ with no additional space requirements, where OPT *stands for an optimal edge (1-)cover of G.*

Sec. 2.3 is dedicated to lifting the assumption on the numerical values. The following definitions are necessary for the discussion of the results we obtain without this assumption:

$$b^{\mathrm{lg}} = \lg\left\lceil \max_{v \in V}\left\{b(v), b(v)^{-1}\right\}\right\rceil \qquad c^{\mathrm{lg}} = \lg\left\lceil \max_{e \in E}\left\{c(e), c(e)^{-1}\right\}\right\rceil \qquad c^{\Delta} = \lg\left\lceil \frac{\max_{e \in E} c(e)}{\min_{e \in E} c(e)}\right\rceil ,$$

where the last parameter captures the number of bits required to encode the edge costs *aspect ratio*.[4] Note that the encoding size $|G|$ of the input weighted hypergraph $G = (V, E, b, c)$ is at least $b^{\mathrm{lg}} + c^{\mathrm{lg}}$. Moreover, c^{Δ} is always at most $2c^{\mathrm{lg}}$, but it may be much smaller than that.

Our results are cast in Thm. 2.2 and 2.3, where the former generalizes Thm. 2.1 and the latter has a better space bound, but slightly worse run-time guarantee. Another drawback of Thm. 2.3 is that it requires that the parameters n and ϵ are known to the algorithm in advance in contrast to Thm. 2.2 and 2.1 that do not require an apriori knowledge of any global parameter.

Theorem 2.2. *On a weighted input hypergraph $G = (V, E, b, c)$, our algorithm uses $O\left(n \log\left(n + m + b^{\mathrm{lg}} + c^{\mathrm{lg}}\right)\right)$ space, processes each input edge $e_t \in E$ in $O(|e_t| \log |e_t|)$ time, and produces a data structure $\mathcal{D}$ with the following guarantee: For every $0 \leq \epsilon < 1$, a $(1 - \epsilon)$-cover certificate χ_ϵ for G such that*

$$c(\mathrm{Im}(\chi_\epsilon)) = O\left(\min\left\{1/\epsilon, \sqrt{n}\right\} \cdot c(\texttt{OPT})\right)$$

can be extracted from $\mathcal{D}$ in time $O(n \log n)$ with no additional space requirements, where OPT *stands for an optimal edge (1-)cover of G.*

Theorem 2.3. *On a weighted input hypergraph $G = (V, E, b, c)$, for any $0 \leq \epsilon < 1$, our algorithm uses $O\left(\log\left(b^{\mathrm{lg}} + c^{\mathrm{lg}}\right) + n \log\left(n + m + c^{\Delta}\right)\right)$ space, processes each input edge $e_t \in E$ in $O(n \log n)$ time, and outputs a $(1 - \epsilon)$-cover certificate χ_ϵ for G such that*

$$c(\mathrm{Im}(\chi_\epsilon)) = O\left(\min\left\{1/\epsilon, \sqrt{n}\right\} \cdot c(\texttt{OPT})\right) ,$$

where OPT *stands for an optimal edge (1-)cover of G.*

[4] Throughout, lg denotes logarithm to the base of 2.

2.1 The Algorithm

In what follows we consider some weighted hypergraph $G = (V, E, b, c)$ with optimal edge (1-)cover OPT. The main building block of algorithm SSSC is a procedure referred to as COVER. This procedure processes the stream of edges and outputs for every node $v \in V$, an identifier of an edge e that covers it, together with an integer variable that intuitively captures the quality of edge e in covering v. Algorithm SSSC uses two parallel invocations of COVER, one on the input graph G and one on some modification of G, and upon termination of the input stream, extracts the desired cover certificate from the output of these two invocations.

2.1.1 Procedure COVER

The procedure maintains for each vertex $v \in V$, the following variables:

- $\text{eid}(v)$ = an identifier $\text{id}(e)$ of some edge $e \in E$; and
- $\text{eff}(v)$ = a (not necessarily positive) integer refereed to as the *effectiveness* of v.

We denote by $\text{eid}_t(v)$ and $\text{eff}_t(v)$ the values of $\text{eid}(v)$ and $\text{eff}(v)$, respectively, at time t (i.e., just before e_t is processed). Procedure COVER that relies on the following definition is presented in Algorithm 1.

Definition (level, effectiveness). Consider edge e_t presented at time t and some subset $T \subseteq e_t$. The *level* of T at time t, denoted $\text{lev}_t(T)$, is defined as

$$\text{lev}_t(T) = \left\lceil \lg \frac{b(T)}{c(e_t)} \right\rceil .$$

Subset T is said to be *effective* at time t if for every $v \in T$, it holds that

$$\text{lev}_t(T) > \text{eff}_t(v) .$$

Note that $\emptyset$ is always vacuously effective.

2.1.2 Algorithm SSSC

We are now ready to present our algorithm SSSC. On input weighted graph $G = (V, E, b, c)$, algorithm SSSC runs in parallel the following procedures that process the stream of edges:

P1: $(\text{eid}_\infty(\cdot), \text{eff}_\infty(\cdot)) \leftarrow$ COVER$(G = (V, E, b, c))$.

P2: $(\text{eid}^{\mathbf{1}}_\infty(\cdot), \text{eff}^{\mathbf{1}}_\infty(\cdot)) \leftarrow$ COVER$(G = (V, E, \mathbf{1}, c))$, where $\mathbf{1}$ stands for the function that assigns a unit benefit to all vertices $v \in V$.

P3: A procedure that maintains for every vertex $v \in V$, a variable $\text{emin}(v)$ that stores the identifier of the minimum cost edge that covers v, seen so far.

Algorithm 1 COVER$(G = (V, E, b, c))$

Initialization $\forall v \in V$: $\mathrm{eid}(v) \leftarrow$ NULL and $\mathrm{eff}(v) \leftarrow -\infty$
for $t = 0, 1, \ldots$ **do**
  Read edge $e_t \in E$ from the stream
  Compute an effective subset $T \subseteq e_t$ of largest benefit $b(T)$
  for all $v \in T$ **do**
    $\mathrm{eid}(v) \leftarrow \mathrm{id}(e_t)$
    $\mathrm{eff}(v) \leftarrow \mathrm{lev}_t(T)$
  end for
end for
return $\mathrm{eid}(\cdot)$ and $\mathrm{eff}(\cdot)$

P4: A procedure that stores for every vertex $v \in V$, its benefit $b(v)$.

Upon termination of the input stream, SSSC takes some parameter $0 \leq \epsilon < 1$ and extracts the desired $(1-\epsilon)$-cover certificate for G from the variables returned by procedures P1–P4. We distinguish between the following two cases.

- Case $\epsilon \geq 1/\sqrt{n}$:
  The algorithm looks for the largest integer r^* such that $b(I(\leq r^*)) \leq \epsilon b(V)$, where

  $$I(\leq r^*) = \{v \in V : \mathrm{eff}_\infty(v) \leq r^*\},$$

  and returns the partial function $\chi : V \to \mathrm{id}(E)$ that maps every vertex $v \in V - I(\leq r^*)$ to $\mathrm{eid}_\infty(v)$.
- Case $\epsilon < 1/\sqrt{n}$:
  The algorithm looks for the largest integer r^* such that $|I^{\mathbf{1}}(\leq r^*)| \leq \sqrt{n}$, where

  $$I^{\mathbf{1}}(\leq r^*) = \{v \in V : \mathrm{eff}^{\mathbf{1}}_\infty(v) \leq r^*\}$$

  and sets χ' to be the partial function $\chi' : V \to \mathrm{id}(E)$ that maps every vertex $v \in V - I^{\mathbf{1}}(\leq r^*)$ to $\mathrm{eid}^{\mathbf{1}}_\infty(v)$. Then, it returns the (complete) function $\chi'' : V \to \mathrm{id}(E)$ extended from χ' by mapping every vertex $v \in I^{\mathbf{1}}(\leq r^*)$ to $\mathrm{emin}(v)$.

Notice that the unweighted case is much simpler: If $G = (V, E)$, then procedure P2 is identical to procedure P1; moreover, procedures P3 and P4 are redundant since all vertices/edges admit a unit benefit/cost. Further note that procedures P1–P4 are oblivious to ϵ. Upon termination of the input stream, the algorithm extracts, for the given $0 \leq \epsilon < 1$, the desired $(1-\epsilon)$-cover certificate for G from the variables returned by procedures P1–P4. In fact, several such cover certificates can be extracted for different values of ϵ.

2.2 Analysis

We begin our analysis with some observations regarding our main procedure COVER.

Observation 2.4. *If $T \subseteq e_t$ is effective at time t and $v \in T$, then $T \cup \{u\}$ is effective at time t for every $u \in e_t$ such that $\mathrm{eff}_t(u) \leq \mathrm{eff}_t(v)$.*

Notice that COVER's updating rule guarantees that the effectiveness $\mathrm{eff}(v)$ is non-decreasing throughout the course of the execution. Employing Obs. 2.4, we can now derive Obs. 2.5 and 2.6 (the former follows by sorting the vertices $v \in e_t$ in non-decreasing order of the value of the effectiveness $\mathrm{eff}(v)$).

Observation 2.5. *The run-time of* COVER *on edge e_t is $O(|e_t| \log |e_t|)$.*

Observation 2.6. *If $T \subseteq e_t$ is effective at time t, then for every $v \in T$, it holds that*

$$\mathrm{eff}_{t+1}(v) \geq \mathrm{lev}_t(T)\,.$$

We are now ready to establish the following lemma.

Lemma 2.7. *Consider some integer r. Procedure* COVER *guarantees that*

$$b\left(\{v \in e_t \mid \mathrm{eff}_{t+1}(v) \leq r\}\right) < 2^{r+1} \cdot c(e_t)\,.$$

Proof. Assume by contradiction that there exists a subset $R \subseteq e_t$, $b(R) \geq 2^{r+1} \cdot c(e_t)$, such that $\mathrm{eff}_{t+1}(v) \leq r$ for every $v \in R$. Since the effectiveness is non-decreasing, it follows that $\mathrm{eff}_t(v) \leq r$ for every $v \in R$, hence the assumption that $b(R) \geq 2^{r+1} \cdot c(e_t)$ ensures that R is effective at time t. But by Obs. 2.6, the effectiveness $\mathrm{eff}_{t+1}(v)$ should have been at least $r+1$ for every $v \in R$, in contradiction to the choice of R. □

Let $\mathrm{eff}_\infty(v)$ denote the value of the variable $\mathrm{eff}(v)$ upon termination of the input stream. Given some integer r, define

$$I(r) = \{v \in V \mid \mathrm{eff}_\infty(v) = r\} \quad \text{and} \quad S(r) = \{e \in E \mid \exists v \in I(r) \text{ s.t. } \mathrm{eid}(v) = \mathrm{id}(e)\}$$

in accordance with the notation defined in Sec. 2.1.2. We extend these two definitions to intervals of integers in the natural way and denote the intervals $(-\infty, r]$ and (r, ∞) in this context by $\leq r$ and $> r$, respectively.

Lemma 2.8. *Consider some integer r. Procedure* COVER *guarantees that*

$$b(I(\leq r)) < 2^{r+1} \cdot c(\mathtt{OPT})\,.$$

Proof. Since the effectiveness is non-decreasing, Lem. 2.7 ensures that for every edge $e \in E$, it holds that

$$b\left(\{v \in e \mid \mathrm{eff}_\infty(v) \leq r\}\right) < 2^{r+1} \cdot c(e)\,.$$

The assertion is established by observing that

$$b(I(\leq r)) \leq \sum_{e \in \texttt{OPT}} b\left(\{v \in e \mid \mathrm{eff}_{\infty}(v) \leq r\}\right) < \sum_{e \in \texttt{OPT}} 2^{r+1} \cdot c(e) = 2^{r+1} \cdot c(\texttt{OPT})\,,$$

where the first inequality is due to the fact that OPT is an edge cover of G. □

Lem. 2.8 will be used to bound from above the benefit of the vertices that are not covered by the edges returned by our algorithm. We now turn to bound from above the cost of these edges.

Lemma 2.9. *Consider some integer r. The edge collection $S(r)$ satisfies*

$$c(S(r)) < b(V)/2^{r-1}\,.$$

Proof. If $e_t \in S(r)$, then there exists some subset $R = R(e_t) \subseteq e_t$ with $\mathrm{lev}_t(R) = r$ such that for every vertex $v \in R$, we have (1) $\mathrm{eff}_t(v) < r$; and (2) $\mathrm{eff}_{t+1}(v) = r$. By definition, the fact that $\mathrm{lev}_t(R) = r$ implies that $c(e_t) < b(R)/2^{r-1}$. Since the variable $\mathrm{eid}(v)$ is updated only when $\mathrm{eff}(v)$ increases and since $\mathrm{eff}(v)$ is non-decreasing, it follows that if $e_t, e_{t'} \in S(r)$, $e_t \neq e_{t'}$, then the subsets $R(e_t)$ and $R(e_{t'})$ are disjoint. Therefore,

$$\sum_{e_t \in S(r)} c(e_t) < \frac{1}{2^{r-1}} \sum_{e_t \in S(r)} b(R(e_t)) \leq b(V)/2^{r-1}$$

which completes the proof. □

The following corollary is obtained by applying Lem. 2.9 to the integers $r+1, r+2, \ldots$

Corollary 2.10. *Consider some integer r. The edge collection $S(> r)$ satisfies*

$$c(S(> r)) < b(V)/2^{r-1}\,.$$

The following important lemma shows that we can extract from the variables returned by COVER an edge subset of low total cost which covers much of the items.

Lemma 2.11. *Consider some $0 < \epsilon < 1$ and let r^* be the largest integer such that $b(I(\leq r^*)) \leq \epsilon \cdot b(V)$. The edge collection $S(> r^*)$ satisfies*

$$c(S(> r^*)) < 8 \cdot c(\texttt{OPT})/\epsilon\,.$$

Proof. Let r be an integer such that $2^{r+1} < \epsilon \cdot \frac{b(V)}{c(\texttt{OPT})} \leq 2^{r+2}$. Lem. 2.8 guarantees that $b(I(\leq r)) < 2^{r+1} \cdot c(\texttt{OPT}) < \epsilon \cdot b(V)$, hence $r \leq r^*$. It follows by Cor. 2.10 that $c(S(> r^*)) \leq c(S(> r)) < b(V)/2^{r-1} \leq 8 \cdot c(\texttt{OPT})/\epsilon$. □

We are now ready to establish the approximation guarantees of algorithm SSSC. Theorem 2.1 (stated under the assumption that all vertex benefits and edge costs are encoded using $O(\log n)$ bits) follows immediately from Theorem 2.12.

Theorem 2.12. *For any $0 \leq \epsilon < 1$, our algorithm outputs a $(1-\epsilon)$-cover certificate for G whose image has cost $O\left(\min\left\{\frac{1}{\epsilon}, \sqrt{n}\right\} \cdot c(\texttt{OPT})\right)$.*

Proof. If $\epsilon \geq 1/\sqrt{n}$, then the assertion follows immediately from Lem. 2.11, so it remains to consider the case of $\epsilon < 1/\sqrt{n}$. We show that χ'' is a 1-cover certificates for G such that $c(\mathrm{Im}(\chi'')) = O(\sqrt{n} \cdot c(\texttt{OPT}))$. Observe first that since $\texttt{OPT}$ covers all vertices in V, it is also an optimal edge 1-cover of $G^{\mathbf{1}}$. Thus, Lem. 2.11 guarantees that $c(\mathrm{Im}(\chi')) < 8\sqrt{n} \cdot c(\texttt{OPT})$. The vertices $v \in V - \mathrm{Dom}(\chi')$ are mapped under χ'' to $\mathrm{emin}(v)$. Since $|V - \mathrm{Dom}(\chi')| \leq \sqrt{n}$ and since $c(\mathrm{emin}(v)) \leq c(\texttt{OPT})$ for every $v \in V$, it follows that

$$c(\mathrm{Im}(\chi'')) < 8\sqrt{n} \cdot c(\texttt{OPT}) + |V - \mathrm{Dom}(\chi')| \cdot c(\texttt{OPT}) \leq 9\sqrt{n} \cdot c(\texttt{OPT}) \,.$$

The assertion follows. □

2.3 Lifting the assumption on the numerical values

We now turn to lift the assumption that all numerical values are encoded using $O(\log n)$ bits and establish Thm. 2.2 and 2.3, starting with the former. To that end, consider the hypergraph $\widetilde{G} = (V, E, \widetilde{b}, \widetilde{c})$ defined by setting $\widetilde{b}(v) = 2^{\lfloor \lg b(v) \rfloor}$ for every vertex $v \in V$ and $\widetilde{c}(e) = 2^{\lfloor \lg c(e) \rfloor}$ for every edge $e \in E$. Since $\widetilde{b}(U)$ and $\widetilde{c}(F)$ are 2-approximations of $b(U)$ and $c(F)$, respectively, for every $U \subseteq V$ and $F \subseteq E$, it follows that a $(1 - O(\epsilon))$-cover certificate for G with image of cost $O\left(\min\left\{\frac{1}{\epsilon}, \sqrt{n}\right\} \cdot c(\texttt{OPT})\right)$ can be obtained by running $\texttt{SSSC}$ on $\widetilde{G}$.

So, in what follows, we assume that $b(v)$ and $c(e)$ are (not necessarily positive) integral powers of 2 for every vertex $v \in V$ and edge $e \in E$. This implies that every benefit $b(v)$ (resp., cost $c(e)$) in G can be encoded using $O(\log b^{\lg})$ (resp., $O(\log c^{\lg})$) bits simply by taking the standard binary representation of $\lg b(v)$ (resp., $\lg c(e)$). Therefore, procedures P3 and P4 can be implemented using $O\left(\log\left(n + m + b^{\lg} + c^{\lg}\right)\right)$ bits per vertex, as desired. Procedure $\texttt{COVER}$ can also be implemented with that many bits per vertex since the level at time t of each subset $T \subseteq e_t$ is an integer whose absolute value satisfies $|\mathrm{lev}_t(T)| = O(b^{\lg} + c^{\lg} + \log n)$, thus establishing Thm. 2.2 due to Obs. 2.5 and Thm. 2.12.

For Thm. 2.3, we need two additional features. First, we scale in an online fashion all vertex benefits and edge costs so that $\min_{v \in V} b(v)$ and $\min_{e \in E} c(e)$ are always 1. We do the same thing with the effectiveness variables $\mathrm{eff}(v)$, only that this time, we ignore those variables with $\mathrm{eff}(v) = -\infty$. This is carried out by maintaining the true values of $\min_{v \in V} b(v)$, $\min_{e \in E} c(e)$, and $\min_{v \in V : \mathrm{eff}(v) > -\infty} \mathrm{eff}(v)$ — denote them by $b_{\min}$, $c_{\min}$, and $\mathrm{eff}_{\min}$, respectively — and scaling all values of $b(v)$, $c(e)$, and $\mathrm{eff}(v)$ stored in the data structures maintained by the procedures of our our algorithm by $b_{\min}$, $c_{\min}$, and $\mathrm{eff}_{\min}$, respectively. Notice that this online scaling requires updating the existing values stored in the data structures whenever $b_{\min}$, $c_{\min}$, or $\mathrm{eff}_{\min}$ are updated, thus resulting in the slightly less favorable run-time promised by Thm. 2.3.

This online scaling feature ensures that the space allocated for the variables of each vertex v is now

$$O\left(\log\left(n+m+b^{\Delta}+c^{\Delta}\right)\right), \tag{1}$$

where $b^{\Delta} = \lg\left\lceil \frac{\max_{v\in V} b(v)}{\min_{v\in V} b(v)} \right\rceil$ is the number of bits required to encode the vertex benefits aspect ratio. We also need additional $O(\log(b^{\lg}+c^{\lg}))$ bits to store the variables $b_{\min}$, $c_{\min}$, and $\text{eff}_{\min}$.

In order to get rid of the dependency on $\log b^{\Delta}$ in (1) and obtain the space bound promised by Thm. 2.3, we use the following feature: Let $\sigma = \sum_{v\in V'} b(v)$, where V' is the set of vertices $v \in V$ encountered by the algorithm so far. Whenever it becomes clear that the contribution of some vertex $v \in V$ to $b(V)$ is at most $\epsilon \cdot b(V)/n$, which is indicated by $b(v) \leq \epsilon\sigma/n$, the algorithm marks vertex v as *insignificant*. Insignificant vertices are treated as if they are not part of the input hypergraph G; in particular, upon marking vertex v as insignificant, the algorithm erases any variable associated with v and updates $b_{\min}$ so that it does not take $b(v)$ into account.

Notice that the total contribution of all insignificant vertices to $b(V)$ is bounded from above by $\epsilon \cdot b(V)$. Therefore, ignoring insignificant vertices cannot hurt our guaranteed coverage by more than an additive term of $\epsilon \cdot b(V)$. The key observation now is that by ignoring insignificant vertices, we keep the parameter b^{Δ} bounded by $b^{\Delta} = O(\log(n/\epsilon))$ as the benefit of any vertex encountered by the algorithm so far is clearly at most σ. Recalling that ϵ is always at least $1/\sqrt{n}$, we conclude that the dependency on $\log b^{\Delta}$ in (1) is replaced by a dependency on $\log\log n$. Thm. 2.3 follows by Thm. 2.12.

3 Lower bounds

A *randomized* semi-streaming algorithm `ALG` for the edge cover problem in hypergraphs is said to be an (n, s, ϵ, ρ)-*algorithm* (resp., an *uncertified* (n, s, ϵ, ρ)-*algorithm*) if given any n-vertex unweighted hypergraph G, `ALG` is guaranteed to maintain a memory of size at most s bits and to output a $(1-\epsilon)$-cover certificate for G with image of expected cardinality at most $\rho \cdot |\texttt{OPT}|$ (resp., to output the identifiers of an edge $(1-\epsilon)$-cover of G whose expected cardinality is at most $\rho \cdot |\texttt{OPT}|$), where `OPT` is an optimal edge cover of G. Our goal in this section is to establish Thm. 3.1 and 3.2, treated in Sec. 3.1 and 3.2, respectively. Observe that the constructions that lie at the heart of Theorems 3.1 and 3.2 are based on hypergraphs whose number of vertices and number of edges are polynomially related, that is, $m = n^{\Theta(1)}$.

Theorem 3.1. *For every integer n_0, there exists an integer $n \geq n_0$ such that for every $\epsilon = \Omega(1/\sqrt{n})$, the existence of an $(n, o(n^{3/2}), \epsilon, \rho)$-algorithm implies that $\rho = \Omega(1/\epsilon)$.*

Theorem 3.2. *Fix some constant real $\alpha > 0$. For every integer n_0, there exists an integer $n \geq n_0$ such that for every $\epsilon \geq n^{-1/2+\alpha}$, the existence of an uncertified $(n, o(n^{1+\alpha}), \epsilon, \rho)$-algorithm implies that $\rho = \Omega\left(\frac{\log\log n}{\log n}\frac{1}{\epsilon}\right)$.*

3.1 The certified case

We shall establish Thm. 3.1 by introducing a probability distribution $\mathcal{G}$ over n-vertex hypergraphs that satisfy the following two properties: (1) Every hypergraph in the support of $\mathcal{G}$ admits an edge cover of cardinality $O(\epsilon\sqrt{n})$. (2) For every *deterministic* semi-streaming algorithm `ALG` that given an n-vertex hypergraph G, maintains a memory of size $o(n^{3/2})$ and outputs a $(1-\epsilon)$-cover certificate χ for G, when `ALG` is invoked on a hypergraph chosen according to $\mathcal{G}$, the expected cardinality of $\mathrm{Im}(\chi)$ is $\Omega(\sqrt{n})$. The theorem than follows by Yao's principle.

3.1.1 The construction of $\mathcal{G}$

Let q be a large prime power. Our construction relies on the *affine plane* $\mathcal{A} = (P, L)$, where P is a set of q^2 *points* and $L \subseteq 2^P$ is a set of $q(q+1)$ *lines* satisfying the following properties:
(1) every line contains q points;
(2) every point is contained in $q+1$ lines;
(3) for every two distinct points, there is exactly one line that contains both of them; and
(4) every two lines intersect in at most one point.
Two lines with an empty intersection are called *parallel*. The line set L can be partitioned into $q+1$ clusters $A_1, \ldots, A_{q+1}$ referred to as *angles*, where $A_i = \{\ell_i^1, \ldots, \ell_i^q\}$ for $i = 1, \ldots, q+1$, such that two distinct lines are parallel if and only if they belong to the same angel. Refer to [16] for an explicit construction of such a combinatorial structure.

Consider some $\frac{1}{3q} \leq \epsilon \leq \frac{1}{66} - \frac{1}{3q}$ and let $r = \lceil 3\epsilon q \rceil$. We construct a random hypergraph $G = (V, E)$ based on the affine plane $\mathcal{A} = (P, L)$ as follows (refer to Figure 1 for an illustration). Fix $V = P$. Randomly partition each line $\ell \in L$ into 2 edges $e_1(\ell) \cup e_2(\ell) = \ell$ by assigning each point in L to one of the 2 edges u.a.r. (and independently of all other random choices).[5] It will be convenient to denote the set of edges corresponding to the lines in angle A_i by $E_i = \{e_1(\ell), e_2(\ell) \mid \ell \in A_i\}$. Let

$$e^* = P - \bigcup_{t=1}^{r} \ell_i^{j(t)} ,$$

where i is an index chosen u.a.r. (and independently) from $[q+1]$ and $1 \leq j(1) < \cdots < j(r) \leq q$ are r distinct indices chosen u.a.r. (and independently) from $[q]$. In other words, e^* is constructed by randomly choosing an angle A_i and then randomly choosing r distinct lines $\ell_i^{j(1)}, \ldots, \ell_i^{j(r)}$ from A_i; the edge consists of all points except those contained in these r lines.

Fix

$$E = E_1 \cup \cdots \cup E_{q+1} \cup \{e^*\} .$$

Observe that $n = |P| = q^2$ and $m = 1 + 2 \cdot |L| = 1 + 2 \cdot q(q+1)$. The execution is divided into two stages, where in the first stage, the edges in $E_1 \cup \cdots \cup E_{q+1}$ are presented in an arbitrary order and

[5] Throughout, we use u.a.r. to abbreviate "uniformly at random".

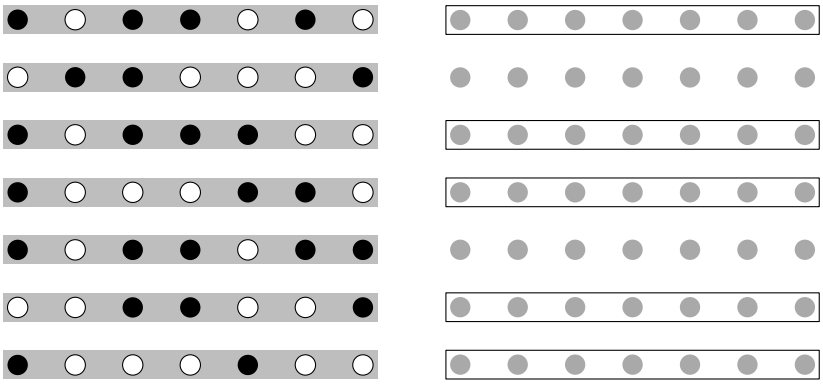

(a) The edges in E_i (b) Edge e^*

Figure 1: The hypergraph G for $q = 7$. (The requirements on ϵ actually imply that q must be larger, but we set $q = 7$ for the sake of a clearer illustration.) The gray rectangles in (a) depict the 7 parallel lines in angle A_i for some $i \in [q+1]$, whereas the black/white circles in each line ℓ_i^j depict the points in $e_1(\ell_i^j)/e_2(\ell_i^j)$. Edge e^*, depicted by the white rectangles in (b), consists of all points except those in $r = 2$ lines of angle A_i.

in the second stage, edge e^* is presented.

3.1.2 Analysis

We start the analysis by observing that G can be covered by the edge e^* and the edges in $\{e_1(\ell_i^{j(t)}), e_2(\ell_i^{j(t)}) \mid 1 \leq t \leq r\}$. Therefore,

$$|\texttt{OPT}| \leq 2r + 1 = O(\epsilon q)\,, \tag{2}$$

where the equation follows from the definition of $r = \lceil 3\epsilon q \rceil$ due to the requirement that $\epsilon \geq \frac{1}{3q}$.

Let s be the space of the deterministic semi-streaming algorithm ALG. Thm. 3.1 is established by combining (2) with the following lemma (that ensures an $\Omega(q)$ expected image cardinality whenever $s = o(n^{3/2})$).

Lemma 3.3. *If $s \leq q^2(q+1)/48$, then w.p. $\geq 1/8$, the $(1-\epsilon)$-cover certificate returned by* ALG *has image of cardinality at least $q/3$.*[6]

Bounding the expected entropy. The proof of Lem. 3.3 is based on information theoretic arguments that require the following definitions. Let X_i^j be a random variable that depicts the partition $(e_1(\ell_i^j), e_2(\ell_i^j))$ of line $\ell_i^j = e_1(\ell_i^j) \cup e_2(\ell_i^j)$ for every $i \in [q+1]$ and $j \in [q]$. Let $X_i = (X_i^1, \ldots, X_i^q)$ and $X = (X_1, \ldots, X_{q+1})$. The independent random choices in the construction of the hypergraph G guarantee that $H(X_i^j) = q$, $H(X_i) = q^2$, and $H(X) = q^2(q+1)$, where $H(\cdot)$

[6] Throughout, we use w.p. and w.h.p. to abbreviate "with probability" and "with high probability", respectively.

denotes the binary entropy function. Before we can proceed with our proof, we have to establish the following lemma whose restriction to the case $k = 1$ is a basic fact in information theory. It will not strike us as a surprise if this lemma was already proved beforehand although we are unaware of any such specific proof; for the sake of completeness, we provide a full proof of this lemma based on Baranyai's Theorem in Appendix A.

Lemma 3.4. *Let $X_1, \ldots, X_n, Y$ be $n+1$ arbitrary random variables and let $1 \leq j(1) < \cdots < j(k) \leq n$ be $1 \leq k \leq n$ distinct indices chosen u.a.r. from $[n]$. Then,*

$$\left\lceil \frac{n}{k} \right\rceil \mathbb{E}_{j(1),\ldots,j(k)} \left[H\left(X_{j(1)}, \ldots, X_{j(k)} \mid Y \right) \right] \;\geq\; H\left(X_1, \ldots, X_n \mid Y \right) .$$

Let M be a random variable that depicts the memory image of `ALG` upon completion of the first stage of the execution. Since M is fully determined by X, it follows that $H(X, M) = H(X)$, hence $H(X \mid M) = H(X) - H(M)$. Recalling that M is described by s bits, we conclude that $H(M) \leq s \leq q^2(q+1)/48$, thus

$$H(X \mid M) \geq \frac{47}{48} \cdot q^2(q+1) = \frac{47}{48} \cdot H(X) . \tag{3}$$

We are now ready to establish the following lemma.

Lemma 3.5. *Our construction guarantees that*

$$\mathbb{P}_{i,j(1),\ldots,j(r)} \left(H\left(X_i^{j(1)}, \ldots, X_i^{j(r)} \mid M \right) \geq \frac{5}{6} \cdot rq \right) \geq 1/4 ,$$

where $i \in [q+1]$ and $1 \leq j(1) < \cdots < j(r) \leq q$ are the random indices chosen during the construction of edge e^*.

Proof. By combining (3) with an application of Lem. 3.4 to the random choice of index $i \in [q+1]$, we derive the inequality

$$\mathbb{E}_i \left[H\left(X_i \mid M \right) \right] \geq \frac{47}{48} \cdot q^2 .$$

Since $H(X_i \mid M) \leq q^2$, we can apply Markov's inequality to conclude that

$$H(X_i \mid M) \geq \frac{23}{24} \cdot q^2 \tag{4}$$

w.p. $\geq 1/2$.

Conditioned on the event that (4) holds, we can apply Lem. 3.4 to the random choice of indices $1 \leq j(1) < \cdots < j(r) \leq q$, deriving the inequality

$$\left\lceil \frac{q}{r} \right\rceil \mathbb{E}_{j(1),\ldots,j(r)} \left[H\left(X_i^{j(1)}, \ldots, X_i^{j(r)} \mid M \right) \right] \geq \frac{23}{24} \cdot q^2$$

which means that

$$\mathbb{E}_{j(1),\ldots,j(r)} \left[H\left(X_i^{j(1)}, \ldots, X_i^{j(r)} \mid M \right) \right] \geq \frac{23}{24} \frac{rq^2}{q+r} .$$

Since $\epsilon \leq \frac{1}{66} - \frac{1}{3q}$, it follows that $r = \lceil 3\epsilon q \rceil \leq 3\epsilon q + 1 \leq q/22$. This, in turn, implies that $\frac{23}{24}\frac{rq^2}{q+r} \geq \frac{11}{12}rq$ which guarantees that

$$\mathbb{E}_{j(1),\ldots,j(r)}\left[H\left(X_i^{j(1)},\ldots,X_i^{j(r)} \mid M\right)\right] \geq \frac{11}{12}\cdot rq\,.$$

Since $H(X_i^{j(1)},\ldots,X_i^{j(r)} \mid M) \leq rq$, we can apply Markov's inequality to conclude that

$$H\left(X_i^{j(1)},\ldots,X_i^{j(r)} \mid M\right) \geq \frac{5}{6}\cdot rq$$

w.p. $\geq 1/2$. The assertion follows as (4) holds w.p. $\geq 1/2$. □

Introducing the random variable Z. Let μ be the actual memory image of ALG upon completion of the first stage of the execution and recall that μ is some instance of the random variable M. Let Z be a real valued random variable that maps the event $M = \mu$ to the entropy in the joint random variable $X_i^{j(1)},\ldots,X_i^{j(r)}$ given $M = \mu$. Observe that by the definition of conditional entropy, we have $\mathbb{E}[Z] = H(X_i^{j(1)},\ldots,X_i^{j(r)} \mid M)$. If the event described in Lem. 3.5 occurs, then $\mathbb{E}[Z] \geq \frac{5}{6}\cdot rq$ and since Z is never larger than rq, we can apply Markov's inequality to conclude that

$$H\left(X_i^{j(1)},\ldots,X_i^{j(r)} \mid M = \mu\right) \geq \frac{2}{3}\cdot rq$$

w.p. $\geq 1/2$. The following corollary is established since the event described in Lem. 3.5 holds w.p. $\geq 1/4$.

Corollary 3.6. *W.p. $\geq 1/8$, the entropy that remains in $X_i^{j(1)},\ldots,X_i^{j(r)}$ after e^* is exposed to* ALG *given that $M = \mu$ is at least $\frac{2}{3}\cdot rq$ bits.*

High entropy implies a large edge cover. Condition hereafter on the event described in Cor. 3.6. Consider the $(1-\epsilon)$-cover certificate χ returned by ALG and let $P' = \bigcup_{t=1}^{r}\ell_i^{j(t)} = P - e^*$ be the set of points not covered by e^*. Let

$$R = \left\{p \in P' \mid p \in \mathrm{Dom}(\chi) \wedge \chi(p) \in E_i\right\}$$

be the set of points not covered by e^* that are mapped under χ to some edge in E_i, where recall that E_i is the set of edges corresponding to the lines in angle A_i (the angle chosen in the random construction of e^*). We can now establish the following lemma.

Lemma 3.7. *Our construction guarantees that $|R| \leq rq/3$.*

Proof. The joint random variable $X_i^{j(1)},\ldots,X_i^{j(r)}$ conditioned on $M = \mu$ can be viewed as a probability distribution π over the matrices $T \in \{1,2\}^{r\times q}$, where $T(t,k) \in \{1,2\}$ indicates whether the k^{th} point in line $\ell_i^{j(t)}$ belongs to edge $e_1(\ell_i^{j(t)})$ or $e_2(\ell_i^{j(t)})$ for every $k \in [q]$ and $1 \leq t \leq r$. Consider some point $p \in R$ and suppose that this is the k^{th} point in line $\ell_i^{j(t)}$. By the definition of

R, all matrices T in the support of π must agree on $T(t,k)$.[7] Therefore, the entropy that remains in $X_i^{j(1)},\ldots,X_i^{j(r)}$ can only arrive from points in $P'-R$. The assertion follows by Cor. 3.6 since each such point contributes at most 1 bit of entropy. □

The cardinality of $\mathrm{Dom}(\chi)$ is at least $|\mathrm{Dom}(\chi)| \geq (1-\epsilon)q^2$. The choice of $r = \lceil 3\epsilon q \rceil$ ensures that $\epsilon q^2 \leq rq/3$, thus $|\mathrm{Dom}(\chi)| \geq q^2 - rq/3$. The key observation now is that even if all these $rq/3$ missing points from $\mathrm{Dom}(\chi)$ are in P', it still leaves us with $|\mathrm{Dom}(\chi) \cap (P'-R)| \geq rq/3$ by Lem. 3.7.

Every point in $\mathrm{Dom}(\chi) \cap (P'-R)$ is covered by some edge $e \in E_j$, $j \neq i$. The properties of the affine plane guarantee that each such edge e covers at most one point in line $\ell_i^{j(t)}$, which sums up to at most r points in P'. Thus, the image of χ must contain (the identifiers of) at least $q/3$ different edges. This concludes the proof of Lem. 3.3. Thm. 3.1 then follows by combining (2) and Lem. 3.3.

3.2 The uncertified case

Similarly to the proof of Thm. 3.1, we shall establish Thm. 3.2 by introducing a probability distribution $\mathcal{G}'$ over n-vertex hypergraphs that this time, satisfies the following two properties: (1) Every hypergraph in the support of $\mathcal{G}'$ admits an edge cover of cardinality $O(\epsilon^2 n)$. (2) For every *deterministic* semi-streaming algorithm `ALG` that given an n-vertex hypergraph $G=(V,E)$, maintains a memory of size $o(n^{1+\alpha})$ and outputs the identifiers of an edge $(1-\epsilon)$-cover $F \subseteq E$ of G, when `ALG` is invoked on a hypergraph chosen according to $\mathcal{G}'$, the expected cardinality of F is $\Omega\left(\epsilon n \frac{\log\log n}{\log n}\right)$. The theorem than follows by Yao's principle.

3.2.1 The construction of $\mathcal{G}'$

We construct a random hypergraph $\hat{G} = (\hat{V}, \hat{E})$ as follows. Let q be a large power of 2 and fix some constant real $\alpha > 0$. Consider some $q^{-(1-\alpha)} \leq \epsilon \leq \frac{1}{66} - \frac{1}{3q}$ and let $r = \lceil 3\epsilon q \rceil$. The main building block of $\hat{G}$ is very similar to the random hypergraph $G=(V,E)$ constructed in Sec. 3.1.1 based on the affine plane $\mathcal{A} = (P,L)$. Specifically, fix $\hat{V} = P$ and let E' be a random edge set constructed just like the construction of the random edge set E presented in Sec. 3.1.1 with the following exception: Instead of randomly partitioning each line $\ell \in L$ into 2 edges $e_1(\ell) \cup e_2(\ell) = \ell$ by assigning each point in L to one of the 2 edges u.a.r. (and independently), we randomly partition each line $\ell \in L$ into r edges $e_1(\ell) \cup \cdots \cup e_r(\ell) = \ell$ by assigning each point in L to one of the r edges u.a.r. (and independently).

[7] In fact, even if we relax the requirement from `ALG` so that χ is allowed to err on some vertices in its domain and the coverage is measured with respect to the vertices for which χ is correct, we can still achieve the desired (asymptotic) bound by using a line of arguments similar to that used in the proof of Lemma 6.2 in [3].

The edge e^* is constructed in the same manner as in Sec. 3.1.1, i.e., we choose an angle A_i u.a.r. and then choose r distinct lines $\ell_i^{j(1)}, \ldots, \ell_i^{j(r)}$ u.a.r. from A_i; the edge consists of all points except those contained in these r lines. (Notice that the parameter r is now used for both the partition of each line into r edges and the construction of edge e^*.) For every $i \in [q+1]$, denote the set of edges corresponding to the lines in angle A_i by $E'_i = \{e_1(\ell), \ldots, e_r(\ell) \mid \ell \in A_i\}$ and fix $E' = E'_1 \cup \cdots \cup E'_{q+1} \cup \{e^*\}$.

The edge multi-set $\hat{E}$ is obtained from E' by augmenting it with *dummy* edges: fix $\hat{E} = E' \cup E_{\mathrm{d}}$, where the edges $e \in E_{\mathrm{d}}$, referred to as dummy edges, are all empty $e = \emptyset$. (Concerns regarding the usage of empty edges can be lifted by augmenting $\hat{V}$ with a dummy vertex v_{d} and taking all dummy edges $e \in E_{\mathrm{d}}$ to be singletons $e = \{v_{\mathrm{d}}\}$.)

Identifier assignment. Recall that the arrival order of the edges is determined by their identifiers so that the edge e_t arriving at time t is assigned with identifier $\mathrm{id}(e_t) = t$. In contrast to the construction presented in Sec. 3.1.1, where the identifier assignment is arbitrary (with the exception that $\mathrm{id}(e^*)$ should be the largest identifier), the assignment of identifiers to the edges in $\hat{E}$ plays a key role in the current construction. Specifically, for every $i \in [q+1]$, $j \in [q]$, and $k \in [r]$, the identifier assigned to edge $e_k(\ell_i^j)$ is

$$\mathrm{id}(e_k(\ell_i^j)) = 0 \circ i \circ j \circ k \circ X_i^{j,k} \,,$$

where i, j, and k are assumed to be encoded as bitstrings of lengths $\lceil \lg(q+1) \rceil$, $\lg q$ (recall that q is a power of 2), and $\lceil \lg r \rceil$, respectively, $\circ$ denotes the string concatenation operator, and $X_i^{j,k}$ is a bitstring of length $3 \lg q$ chosen u.a.r. (and independently). Notice that each identifier contains $\iota = 1 + \lceil \lg(q+1) \rceil + \lg q + \lceil \lg r \rceil + 3 \lg q$ bits encoding some integer (with the most significant bit on the left) in $[0, 2^{\iota-1} - 1]$ and by design, each edge in $E'_1 \cup \cdots \cup E'_{q+1}$ is assigned with a unique identifier.

The identifier assigned to edge e^* is $\mathrm{id}(e^*) = 1 \circ 0^{\iota-1}$, which encodes the integer $2^{\iota-1}$. The dummy edges are used for filling up the gaps between the identifiers assigned to the edges in E' so that $\mathrm{id}(\cdot)$ is a bijection from $\hat{E} = E' \cup E_{\mathrm{d}}$ to $\left[0, 2^{\iota-1}\right]$. As e^* is assigned with the highest identifier, this is the last edge to arrive. Observe that $n = q^2$ and $m = 2^{\iota-1} + 1 = O(q^6)$.

3.2.2 Analysis

We start the analysis by observing that $\hat{G}$ can be covered by edge e^* and the edges in $\{e_1(\ell_i^{j(t)}), \ldots, e_r(\ell_i^{j(t)}) \mid 1 \leq t \leq r\}$. Therefore,

$$|\texttt{OPT}| \leq r^2 + 1 = O(\epsilon^2 q^2) \,, \tag{5}$$

where the equation follows from the definition of $r = \lceil 3\epsilon q \rceil$ due to the requirement that $\epsilon = \omega(q^{-1})$.

Let s be the space of the deterministic semi-streaming algorithm ALG. Thm. 3.2 is established by combining (5) with the following lemma (that ensures an $\tilde{\Omega}(\epsilon q^2)$ expected set cover cardinality whenever $s = o(\epsilon n^{3/2})$).

Lemma 3.8. *If $s \leq rq(q+1)/16$, then w.p. $\geq 1/9$, the edge $(1-\epsilon)$-cover returned by* ALG *has cardinality $\Omega\left(\epsilon q^2 \frac{\log\log q}{\log q}\right)$.*

The proof of Lem. 3.8 is based on information theoretic arguments that require the following definitions. Recall that $X_i^{j,k}$ is a random bitstring of length $3\lg q$ used in the construction of $\mathrm{id}(e_k(\ell_i^j))$ for every $i \in [q+1]$, $j \in [q]$, and $k \in [r]$. Let $X_i^j = (X_i^{j,1}, \ldots, X_i^{j,r})$, $X_i = (X_i^1, \ldots, X_i^q)$, and $X = (X_1, \ldots, X_{q+1})$. The independent random choices in the construction of the identifiers of $\hat{E}$ guarantee that $H(X_i^{j,k}) = 3\lg q$, $H(X_i^j) = 3r\lg q$, $H(X_i) = 3rq\lg q$, and $H(X) = 3rq(q+1)\lg q$.

As in the analysis performed in Sec. 3.1.2, let $i \in [q+1]$ and $1 \leq j(1) < \cdots < j(r) \leq q$ be the random indices chosen in the construction of edge e^*. Let M be a random variable that depicts the memory image of ALG before the last edge e^* arrives and let μ be its actual instantiation. Observing that $H(X \mid M) \geq \frac{47}{48} \cdot H(X)$ (cf. inequality (3)), we can repeat the line of arguments used in Sec. 3.1.2 to derive the following corollary (analogous to Cor. 3.6).

Corollary 3.9. *W.p. $\geq 1/8$, the entropy that remains in $X_i^{j(1)}, \ldots, X_i^{j(r)}$ after e^* is exposed to* ALG *given that $M = \mu$ is at least $2r^2 \lg q$ bits.*

Notice that the requirement $\epsilon \geq q^{-(1-\alpha)}$ ensures that $r = \lceil 3\epsilon q \rceil$ and q are polynomially related and so are r and $n = q^2 + 1$. Therefore, an event that holds w.h.p. with respect to the parameter r also holds w.h.p. with respect to the parameters q and n; in what follows, whenever we use the term w.h.p., we refer to w.h.p. with respect to these three parameters.

Lemma 3.10. *W.h.p., all edges $e_k(\ell_i^{j(t)})$, $t \in [r], k \in [r]$, satisfy $(5/6)q/r \leq |e_k(\ell_i^{j(t)})| \leq 2q/r$.*

Proof. Fix some $t \in [r]$ and $k \in [r]$. The random partition of line $\ell_i^{j(t)}$ into the r edges $e_1(\ell_i^{j(t)}) \cup \cdots \cup e_r(\ell_i^{j(t)}) = \ell_i^{j(t)}$ implies that $\mathbb{E}[|e_k(\ell_i^{j(t)})|] = q/r$. By Chernoff's bound, we have $(5/6)q/r \leq |e_k(\ell_i^{j(t)})| \leq 2q/r$ w.h.p. The assertion follows by union bound. □

Identifiers with large entropy. Condition hereafter on the events described in Cor. 3.9 and Lem. 3.10. Since Cor. 3.9 ensures that

$$\sum_{t=1}^{r}\sum_{k=1}^{r} H(X_i^{j(t),k} \mid M = \mu) \geq H(X_i^{j(1)}, \ldots, X_i^{j(r)} \mid M = \mu) \geq 2r^2 \lg q$$

and since $H(X_i^{j(t),k} \mid M = \mu) \leq 3\lg q$ for every $(t,k) \in [r] \times [r]$, it follows that there exists a subset $\Psi \subseteq [r] \times [r]$ such that (1) $|\Psi| \geq r^2/2$; and (2) $H(X_i^{j(t),k} \mid M = \mu) \geq \lg q$ for every $(t,k) \in \Psi$.

Consider some pair $(t,k) \in \Psi$. The definition of Ψ guarantees that at least $\lg q$ bits of entropy remain in the identifier $\mathrm{id}(e_k(\ell_i^{j(t)}))$ of edge $e_k(\ell_i^{j(t)})$ after e^* is exposed to ALG given that $M = \mu$. Thus, ALG must have at least q different candidates for $\mathrm{id}(e_k(\ell_i^{j(t)}))$. The design of the identifier

assignment function $\mathrm{id}(\cdot)$ guarantees that all but one of these candidate identifiers are actually assigned to dummy edges and that the candidate identifiers of edge $e_k(\ell_i^{j(t)})$ and the candidate identifiers of edge $e_{k'}(\ell_i^{j(t')})$ are disjoint for every $(t,k),(t',k') \in \Psi$, $(t,k) \neq (t',k')$. Therefore, every edge $e_k(\ell_i^{j(t)})$ with $(t,k) \in \Psi$ that is guaranteed to belong to the edge $(1-\epsilon)$-cover F output by `ALG` contributes at least q distinct edges to $|F|$.

On the other hand, Lem. 3.10 ensures that the points in $e_k(\ell_i^{j(t)})$ can be covered by at most $2q/r \ll q$ edges belonging to $E'_{-i} = E'_1 \cup \cdots \cup E'_{i-1} \cup E'_{i+1} \cup \cdots \cup E'_{q+1}$, that is, edges corresponding to lines of angles other than A_i. Hence, for the sake of the analysis, we may assume hereafter that `ALG` covers the points in $e_k(\ell_i^{j(t)})$ by edges belonging to E'_{-i} for every $(t,k) \in \Psi$.

Coverage from another angle. Let $N = \bigcup_{(t,k)\in\Psi} e_k(\ell_i^{j(t)})$ be the set of points contained in the edges corresponding to the index pairs in Ψ. Since $|\Psi| \geq r^2/2$ and since Lem. 3.10 guarantees that $|e_k(\ell_i^{j(t)})| \geq (5/6)q/r$ for every $(t,k) \in \Psi$, it follows that $|N| \geq 5qr/12$.

Recall that the edge $(1-\epsilon)$-cover F may leave at most ϵq^2 uncovered points. The choice of $r = \lceil 3\epsilon q \rceil$ ensures that $\epsilon q^2 \leq qr/3$, thus at most $qr/3$ points are not covered by F. The key observation now is that even if all these uncovered points belong to N, then F should still cover at least $5qr/12 - qr/3 = qr/12$ points in N; let $N' \subseteq N$ be the subset consisting of these (at least) $qr/12$ covered points.

We argue that in order to cover the points in N' with edges belonging to E_{-i}, one needs $\Omega\left(\epsilon q^2 \frac{\log\log q}{\log q}\right) = \Omega\left(qr \frac{\log\log q}{\log q}\right)$ distinct edges w.h.p. The proof of Lem. 3.8 is completed by union bound since the events described in Cor. 3.9 and Lem. 3.10 (i.e., the events on which our analysis is conditioned) hold w.p. $\geq 1/8$ and w.h.p., respectively. To that end, consider some line $\ell \in L - A_i$, namely, a line from an angle other than A_i. The properties of the affine plane $\mathcal{A}$ ensure that the intersection $I(\ell) = \ell \cap (\ell_i^{j(1)} \cup \cdots \cup \ell_i^{j(r)})$ contains exactly $|I(\ell)| = r$ points. The assignment of these r points to the edges $e_1(\ell), \ldots, e_r(\ell)$ is determined by the random partition of ℓ into $e_1(\ell) \cup \cdots \cup e_r(\ell) = \ell$ and it can be viewed as a balls-into-bins process with r balls and r bins. By a known result on balls-into-bins processes (see, e.g., [18]), we conclude that w.h.p., $\max_{k\in[r]} |e_k(\ell) \cap I(\ell)| = O\left(\frac{\log r}{\log\log r}\right)$ and by union bound, this holds for all lines $\ell \in L - A_i$ w.h.p.; in particular, every edge in E'_{-i} covers $O\left(\frac{\log r}{\log\log r}\right)$ points in N' . The argument follows since $|N'| = \Omega(qr)$.

This concludes the proof of Lem. 3.8. Thm. 3.2 then follows by combining (5) and Lem. 3.8.

APPENDIX

A Proving Lem. 3.4

Assume first that $n/k = d$ for some integer $d \geq 1$. Let $\mathcal{S}(n,k)$ be the collection of all $\binom{n}{k}$ subsets $S \subseteq [n]$ of cardinality $|S| = k$. By Baranyai's Theorem (see, e.g., [20]), there exists a partition $\mathcal{P}$ of $\mathcal{S}(n,k)$ into $\binom{n}{k}/d$ pairwise disjoint clusters such that every cluster C of $\mathcal{P}$ consists of d subsets $S \in \mathcal{S}(n,k)$ whose union satisfies $\bigcup_{S \in C} S = [n]$. Note that by definition, the subsets in C must be pairwise disjoint.

Given some subset $S = \{j_1, \ldots, j_\ell\} \subseteq [n]$, let X_S denote the joint random variable $(X_{j_1}, \ldots, X_{j_\ell})$. Fix some cluster $C = \{S_1, \ldots, S_d\}$ of $\mathcal{P}$. The chain rule of conditional entropy implies that

$$\begin{aligned} H\left(X_1, \ldots, X_n \mid Y\right) &= H\left(X_{S_1} \mid Y\right) + H\left(X_{S_2} \mid X_{S_1} \mid Y\right) + \cdots + H\left(X_{S_d} \mid X_{S_1 \cup \cdots \cup S_{d-1}} \mid Y\right) \\ &\leq H\left(X_{S_1} \mid Y\right) + H\left(X_{S_2} \mid Y\right) + \cdots + H\left(X_{S_d} \mid Y\right) . \end{aligned}$$

Denoting the clusters of $\mathcal{P}$ by $C^1, \ldots, C^{\binom{n}{k}/d}$ and letting $C^i = \{S_1^i, \ldots, S_d^i\}$ for $i = 1, \ldots, \binom{n}{k}/d$, we can sum over all clusters of $\mathcal{P}$ to conclude that

$$\frac{\binom{n}{k}}{d} H\left(X_1, \ldots, X_n \mid Y\right) \leq \sum_{i=1}^{\binom{n}{k}/d} \sum_{j=1}^{d} H\left(X_{S_j^i} \mid Y\right) . \tag{A-1}$$

The assertion follows since the right hand side of (A-1) has $\binom{n}{k}$ terms, each identified with a unique subset $S \in \mathcal{S}(n,k)$, hence if we pick one term u.a.r., then its expected value is at least $H(X_1, \ldots, X_n \mid Y)/d$.

Now, assume that $n = k \cdot d - r$ for some integers $d \geq 1$ and $0 < r < k$ and let $n' = k \cdot d$. Let $X_{n+1}, \ldots, X_{n'}$ be r *dummy* random variables with 0 entropy. We have all ready showed that if subset $S \subseteq [n']$ is chosen u.a.r. from $\mathcal{S}(n',k)$, then

$$d \cdot \mathbb{E}_S\left[H\left(X_S \mid Y\right)\right] \geq H\left(X_1, \ldots, X_{n'} \mid Y\right) = H\left(X_1, \ldots, X_n \mid Y\right) .$$

Since $H\left(X_S \mid Y\right) = H\left(X_{S \cap [n]} \mid Y\right)$ for every $S \in \mathcal{S}(n',k)$, it follows that shifting the probability mass in a uniform manner from subsets S containing dummy variables to subsets S that do not contain dummy variables cannot decrease the expected entropy; in other words, if subset $S \subseteq [n]$ is chosen u.a.r. from $\mathcal{S}(n,k)$ and subset $S' \subseteq [n']$ is chosen u.a.r. from $\mathcal{S}(n',k)$, then

$$\mathbb{E}_S\left[H\left(X_S \mid Y\right)\right] \geq \mathbb{E}_{S'}\left[H\left(X_S \mid Y\right)\right] .$$

The assertion follows since $d = \lceil n/k \rceil$.

References

[1] K. Ahn and S. Guha. Graph sparsification in the semi-streaming model. In *ICALP*, pages 328–338, 2009.

[2] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. The online set cover problem. *SIAM J. Comput.*, 39(2):361–370, 2009.

[3] N. Alon, Y. Emek, M. Feldman, and M. Tennenholtz. Adversarial leakage in games. *SIAM J. Discrete Math.*, 27(1):363–385, 2013.

[4] Y. Emek, M. M. Halldórsson, and A. Rosén. Space-constrained interval selection. In *ICALP (1)*, pages 302–313, 2012.

[5] L. Epstein, A. Levin, J. Mestre, and D. Segev. Improved approximation guarantees for weighted matching in the semi-streaming model. In *STACS*, pages 347–358, 2010.

[6] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. On graph problems in a semi-streaming model. *Theor. Comput. Sci.*, 348:207–216, 2005.

[7] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. Graph distances in the data-stream model. *SIAM J. Comput.*, 38(5):1709–1727, 2008.

[8] P. Fraigniaud, M. M. Halldórsson, B. Patt-Shamir, D. Rawitz, and A. Rosén. Shrinking maxima, decreasing costs: New online packing and covering problems. In *APPROX-RANDOM*, pages 158–172, 2013.

[9] F. Grandoni, A. Gupta, S. Leonardi, P. Miettinen, P. Sankowski, and M. Singh. Set covering with our eyes closed. *SIAM J. Comput.*, 42(3):808–830, 2013.

[10] B. V. Halldórsson, M. M. Halldórsson, E. Losievskaja, and M. Szegedy. Streaming algorithms for independent sets. In *ICALP*, pages 641–652, 2010.

[11] L. Jia, G. Lin, G. Noubir, R. Rajaraman, and R. Sundaram. Universal approximations for tsp, steiner tree, and set cover. In *STOC*, pages 386–395, 2005.

[12] R. M. Karp. Reducibility Among Combinatorial Problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[13] J. A. Kelner and A. Levin. Spectral sparsification in the semi-streaming setting. *Theory Comput. Syst.*, 53(2):243–262, 2013.

[14] C. Konrad, F. Magniez, and C. Mathieu. Maximum matching in semi-streaming with few passes. In *APPROX*, pages 231–242, 2012.

[15] I. Kremer, N. Nisan, and D. Ron. On randomized one-round communication complexity. *Computational Complexity*, 8(1):21–49, 1999.

[16] C. C. Lindner and C. A. Rodger. *Design Theory*. Discrete Mathematics and its Applications. CRC Press, 2nd edition, 2011.

[17] A. McGregor. Finding graph matchings in data streams. In *APPROX-RANDOM*, pages 170–181, 2005.

[18] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Jan. 2005.

[19] B. Saha and L. Getoor. On maximum coverage in the streaming model & application to multi-topic blog-watch. In *SDM*, pages 697–708, 2009.

[20] J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, 2nd edition, 2001.

[21] V. V. Vazirani. *Approximation algorithms*. Springer-Verlag New York, Inc., New York, NY, USA, 2001.